

Nearly all the sky is covered by Lyman- α emission around high-redshift galaxies

L. Wisotzki^{1*}, R. Bacon², J. Brinchmann^{3,4}, S. Cantalupo⁵, P. Richter⁶, J. Schaye³, K. B. Schmidt¹, T. Urrutia¹, P. M. Weillbacher¹, M. Akhlaghi², N. Bouché⁷, T. Contini⁷, B. Guiderdoni², E. C. Herenz⁸, H. Inami², J. Kerutt¹, F. Leclercq², R. A. Marino⁵, M. Maseda³, A. Monreal-Ibero^{9,10}, T. Nanayakkara³, J. Richard², R. Saust¹, M. Steinmetz¹ & M. Wendt^{1,6}

Galaxies are surrounded by large reservoirs of gas, mostly hydrogen, that are fed by inflows from the intergalactic medium and by outflows from galactic winds. Absorption-line measurements along the lines of sight to bright and rare background quasars indicate that this circumgalactic medium extends far beyond the starlight seen in galaxies, but very little is known about its spatial distribution. The Lyman- α transition of atomic hydrogen at a wavelength of 121.6 nanometres is an important tracer of warm (about 10^4 kelvin) gas in and around galaxies, especially at cosmological redshifts greater than about 1.6 at which the spectral line becomes observable from the ground. Tracing cosmic hydrogen through its Lyman- α emission has been a long-standing goal of observational astrophysics^{1–3}, but the extremely low surface brightness of the spatially extended emission is a formidable obstacle. A new window into circumgalactic environments was recently opened by the discovery of ubiquitous extended Lyman- α emission from hydrogen around high-redshift galaxies^{4,5}. Such measurements were previously limited to especially favourable systems^{6–8} or to the use of massive statistical averaging^{9,10} because of the faintness of this emission. Here we report observations of low-surface-brightness Lyman- α emission surrounding faint galaxies at redshifts between 3 and 6. We find that the projected sky coverage approaches 100 per cent. The corresponding rate of incidence (the mean number of Lyman- α emitters penetrated by any arbitrary line of sight) is well above unity and similar to the incidence rate of high-column-density absorbers frequently detected in the spectra of distant quasars^{11–14}. This similarity suggests that most circumgalactic atomic hydrogen at these redshifts has now been detected in emission.

We used the Multi-Unit Spectroscopic Explorer (MUSE), developed by our team and installed in 2014 at the ESO Very Large Telescope^{15–17}, to perform very long exposures in two fields that were previously mapped to extreme depths with the Hubble Space Telescope (HST): the Hubble Deep Field South¹⁸ (HDFS) and the Hubble Ultra Deep Field¹⁹ (HUDF). In our MUSE data we detected 270 Lyman- α ($\text{Ly}\alpha$)-emitting galaxies at $3 < z < 6$ (Methods), many of which are barely visible or even undetected with the HST. Extended $\text{Ly}\alpha$ haloes around these galaxies can be traced to distances of a few arcseconds from the source centres^{4,5}.

In a first approach to estimating the total sky coverage of extended $\text{Ly}\alpha$ emission, we constructed redshift-integrated $\text{Ly}\alpha$ maps in the two observed fields as follows (Methods): we extracted pseudo-narrowband $\text{Ly}\alpha$ subimages around each selected object from the MUSE data. Coadding all subimages over the full redshift range followed by some spatial filtering yielded a $\text{Ly}\alpha$ image for each field. Figure 1 shows this image for the HUDF. Counting the number of pixels above a given $\text{Ly}\alpha$ surface brightness $s_{\text{Ly}\alpha}$ yields the fractional sky coverage $f_{\text{Ly}\alpha}$. For a threshold of $s_{\text{Ly}\alpha} > 10^{-19} \text{ erg s}^{-1} \text{ cm}^{-2} \text{ arcsec}^{-2}$ —the typical 1σ limiting

surface brightness in the narrowband images inside an aperture of $1''$ —we find a sky coverage of 46% in the HUDF and 45% in the HDFS (Methods). Although this result already suggests that the sky coverage might further increase for even lower thresholds, the approach is hampered by noise and the need to apply spatial filtering.

To lower the surface brightness limit beyond the sensitivity of individual lines of sight, we employed a combination of azimuthal averaging and image stacking (Methods). We first computed radial $\text{Ly}\alpha$ surface brightness profiles for each object by averaging over pre-defined concentric annuli. Motivated by the fact that our $\text{Ly}\alpha$ halo profiles do not, on average, depend strongly on $\text{Ly}\alpha$ luminosity⁵, we median-combined the individual images in three redshift bins (Methods) and approximated the radial profiles of the median images by smooth fitting functions. Figure 2 illustrates this process. In the outermost annuli,




Fig. 1 | Distribution of the observed $\text{Ly}\alpha$ emission in the HUDF. The underlying image is a colour composite obtained by the HST¹⁹ restricted to the $1' \times 1'$ section observed with MUSE. The extended $\text{Ly}\alpha$ emission detected by MUSE is superimposed in blue, summed over the redshift range $3 < z < 6$ and spatially filtered to suppress the noise. The grey semi-transparent areas outline the MUSE field of view and also mask the brightest foreground galaxies. The dynamic range of the $\text{Ly}\alpha$ overlay was adjusted such that the faintest visible structures have a surface brightness of $10^{-19} \text{ erg s}^{-1} \text{ cm}^{-2} \text{ arcsec}^{-2}$. Credit: NASA, ESA, and S. Beckwith (STScI) and the HUDF Team (CC BY 4.0; <https://creativecommons.org/licenses/by/4.0/>); crop, alignment, and application of blur filter by L.W.

¹Leibniz-Institut für Astrophysik Potsdam (AIP), Potsdam, Germany. ²Université Lyon, Université Lyon 1, ENS de Lyon, CNRS, Centre de Recherche Astrophysique de Lyon, Saint-Genis-Laval, France. ³Leiden Observatory, Leiden University, Leiden, The Netherlands. ⁴Instituto de Astrofísica e Ciências do Espaço, Universidade do Porto, CAUP, Porto, Portugal. ⁵Department of Physics, ETH Zürich, Zürich, Switzerland. ⁶Institut für Physik und Astronomie, Universität Potsdam, Potsdam, Germany. ⁷Institut de Recherche en Astrophysique et Planétologie (IRAP), Université de Toulouse, CNRS, Toulouse, France. ⁸Department of Astronomy, Stockholm University, AlbaNova University Centre, Stockholm, Sweden. ⁹Instituto de Astrofísica de Canarias (IAC), La Laguna, Tenerife, Spain. ¹⁰Universidad de La Laguna, Departamento Astrofísica, La Laguna, Tenerife, Spain. *e-mail: lwisotzki@aip.de

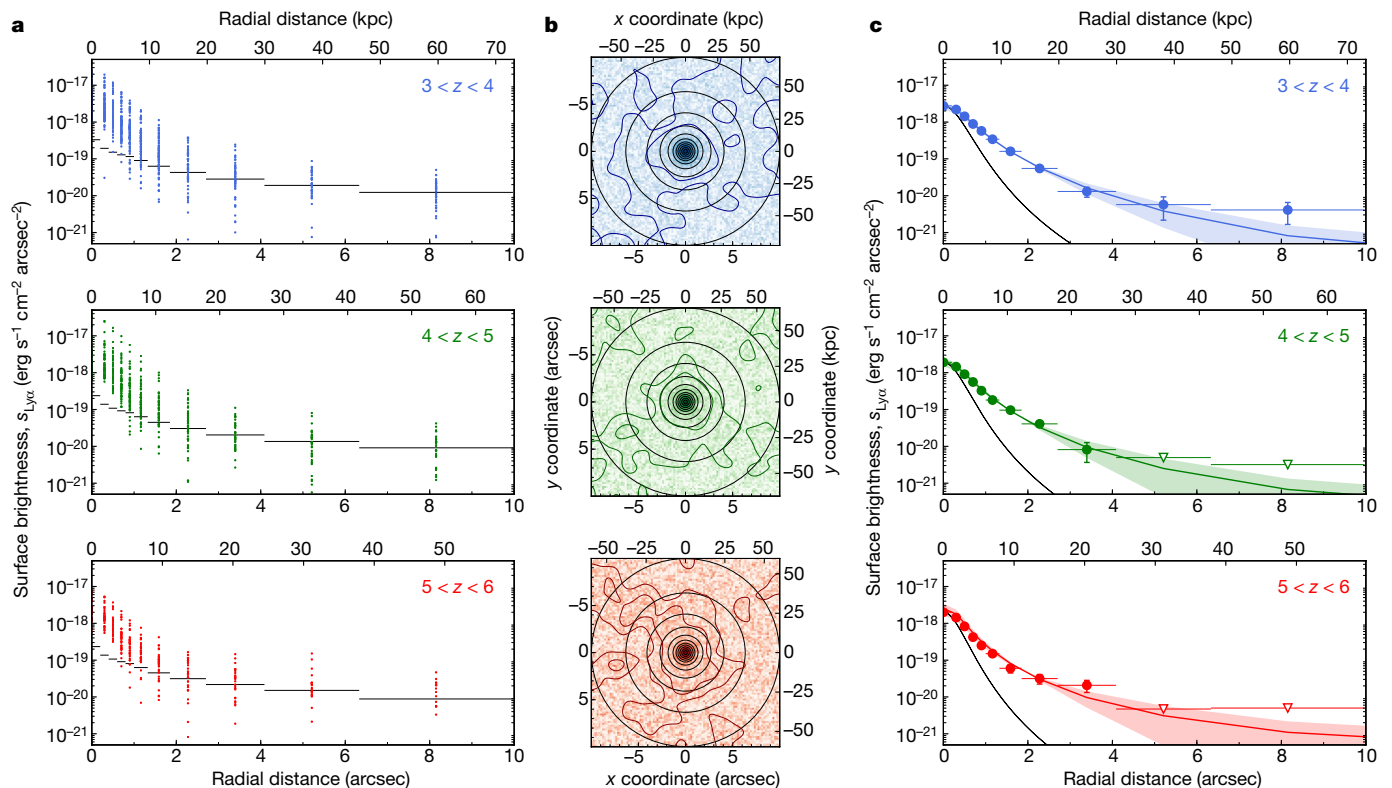


Fig. 2 | Stacking and construction of representative radial Ly α profiles.

a, Ly α surface brightness measurements of the individual galaxies in the HUDF and HDFS, azimuthally averaged in concentric annuli (dots). The surface brightness levels indicated by the horizontal error bars correspond to the typical 1σ uncertainties, and their horizontal extents give the widths of the annuli. The three panels (top, middle and bottom) represent three disjoint redshift ranges as given by the labels at top right of the panels. Radial coordinates are given in angular (bottom axes) as well as physical (top axes) units, the latter evaluated at the centre of each redshift range assuming a cosmological model with $h = 0.7$, $\Omega_m = 0.3$ and $\Omega_\Lambda = 0.7$. **b**, Median-stacked Ly α images for these redshift bins (indicated by colours to match those in **a**). The contours trace surface brightnesses of $(0.5, 2) \times 10^{-20} \text{ erg s}^{-1} \text{cm}^{-2} \text{arcsec}^{-2}$ after subtracting a model image,

the median-stacked profiles reach limiting surface brightness levels of $s_{\text{Ly}\alpha, \text{lim}} \approx (5, 4, 4) \times 10^{-21} \text{ erg s}^{-1} \text{cm}^{-2} \text{arcsec}^{-2}$ (1σ , at $z \approx 3.5, 4.5, 5.5$), an order of magnitude more sensitive than for the single line of sight measurements considered above.

From the scaled median-stacked profiles we created synthetic Ly α maps for the three redshift bins and for the full redshift range. These maps, shown in Fig. 3a, represent idealized, noise-free and approximately seeing-corrected models of the Ly α distribution in the sky. The resulting cumulative fractional Ly α sky coverage $f_{\text{Ly}\alpha}$ is presented in Fig. 3b as a function of the surface brightness threshold. At $s_{\text{Ly}\alpha} \approx 10^{-20} \text{ erg s}^{-1} \text{cm}^{-2} \text{arcsec}^{-2}$, $f_{\text{Ly}\alpha}$ is already well above 80% and still increasing. At these extremely faint levels, the contributions of $f_{\text{Ly}\alpha}$ from the different redshift ranges formally add up to more than 100%, a clear sign that the Ly α emission regions substantially overlap in projection.

The sky coverage is an intuitively appealing number but of limited use as it saturates at 100%. A closely related but more physically useful quantity is the incidence rate dn/dz , the average number of Ly α -emitting regions per unit redshift passed by a typical line of sight, at a given surface brightness level. This quantity can be directly compared to dn/dz (obtained from absorption line statistics) for different absorber column densities. We also corrected for cosmological surface brightness dimming by moving from observed surface brightness $s_{\text{Ly}\alpha}$ to intrinsic ‘surface luminosity’ $S_{\text{Ly}\alpha}$ expressed in $\text{erg s}^{-1} \text{kpc}^{-2}$. Furthermore, we accounted for the inevitable faint-end incompleteness of the Ly α

smoothing the residual with a Gaussian of $2''$ full-width at half-maximum (FWHM) and adding back the model. The overplotted black circles show the boundaries of the annuli used to extract the radial profiles.

c, Azimuthally averaged radial profiles of the median-stacked images (data points), with inverted triangles indicating upper limits. The vertical bars on the data points quantify the estimated 1σ errors, while the horizontal bars again indicate the widths of the annuli. The black line in each panel traces the radial shape of a scaled point source^{4,5}, demonstrating that the median Ly α emission is well resolved for radii greater than about 1 arcsec. The solid, coloured curves show the profiles extracted from two-dimensional surface brightness model fits to the median images, with the shaded regions indicating the estimated 1σ uncertainties.

emitter sample by tying the integration to a completeness-corrected population distribution statistic (Methods). The resulting cumulative incidence rates as functions of surface luminosity threshold are presented in Fig. 4a. Values of $dn/dz > 1$ indicate that a random line of sight passes on average through more than one emitter within a redshift interval of $\Delta z = 1$.

In Fig. 4b we compare our measured incidence rates with the statistics of atomic hydrogen detected in absorption against background quasars^{11–14}. We find that emission and absorption incidence rates dn/dz_{em} and dn/dz_{abs} have a similar range of values, which we use to tentatively match surface luminosities to column densities.

Emission regions with $\log_{10}[S_{\text{Ly}\alpha} (\text{erg s}^{-1} \text{kpc}^{-2})] \gtrsim 38$ (for brevity we omit the units of $S_{\text{Ly}\alpha}$ in the following), typically at radial distances of less than about $2''$, have a dn/dz_{em} of about 0.5 per unit redshift, which is comparable to damped Ly α absorbers^{11,14} (DLAs) with column densities of $\log_{10}[N(\text{H I}) (\text{cm}^{-2})] > 20.3$. At redshifts $z \lesssim 3.5$ this result broadly agrees with previous findings²⁰ based on long-slit spectroscopy of a much smaller sample. Our data also show that the trend of dn/dz with redshift is very similar for absorbers and emitters. It is thus plausible to identify DLAs with Ly α -emitting regions at levels of $\log_{10}(S_{\text{Ly}\alpha}) > 38$, which is also approximately the limit for the detection of individual Ly α haloes^{4,5}. This Ly α emission is likely to be powered by ultraviolet photons from star-forming regions and then resonantly scattered outwards^{21–23}, possibly enhanced by cooling radiation during the accretion of gas into dark matter haloes^{24–26}.

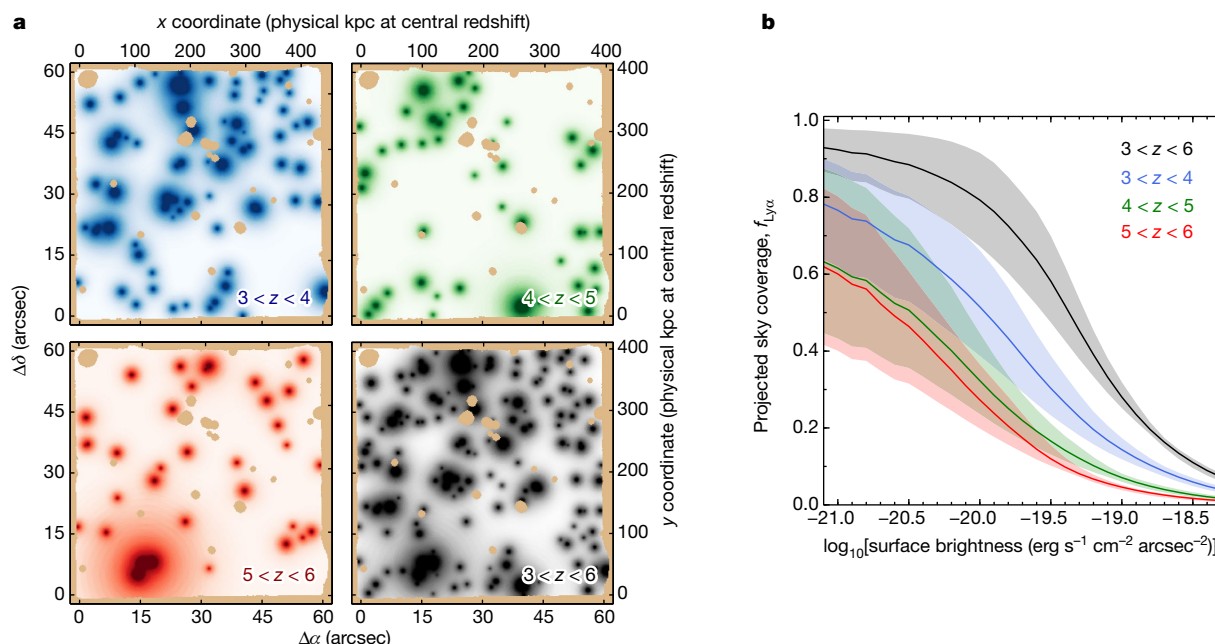


Fig. 3 | The Ly α sky coverage from median-stacked profiles.

a, Reconstructed noise-free Ly α model images of the $1' \times 1'$ section of the HUDF, shown separately for the three disjoint redshift bins (given at bottom right of each panel, and indicated by the blue/green/red colours) and for the full redshift range (shown in black). The shading indicates surface brightness in logarithmic stretch. At the observed position of each object a source model was inserted after rescaling it to its actual Ly α flux. The

light brown areas delineate the MUSE field of view and also mask bright foreground objects. **b**, The cumulative fractional sky coverage of projected Ly α emission, as a function of limiting surface brightness. Labels and colours indicate the redshift ranges. The solid lines represent the average relations from our two MUSE pointings, while the shaded bands outline our 1σ uncertainty estimates from combining the errors of the profiles and the differences between the two fields.

Moving to lower column densities of atomic hydrogen, Fig. 4b shows that systems with $\log_{10}[N(\text{H I}) (\text{cm}^{-2})] > 19$ (sometimes called sub-DLAs or SDLAs) and Ly α emission regions with $\log_{10}(S_{\text{Ly}\alpha}) \gtrsim 37.5$ both have dn/dz of order unity. At even lower $N(\text{H I})$, Lyman limit systems (LLSs) with $\log_{10}[N(\text{H I}) (\text{cm}^{-2})] > 17$ give rise to several incidences per unit redshift¹², which can be approximately matched in emission

by surface luminosities of $\log_{10}(S_{\text{Ly}\alpha}) \gtrsim 37$, a level just detectable in our median stack at $3 < z < 4$ (but requiring some extrapolation of the profiles for $z > 4$). Above column densities of about 10^{18} cm^{-2} , hydrogen becomes self-shielded to ionizing radiation and thus at least partly atomic²⁷, although most of the gas in this $N(\text{H I})$ range is still ionized. In fact nearly all atomic hydrogen at $z > 3.5$ is found¹¹ in absorbers with

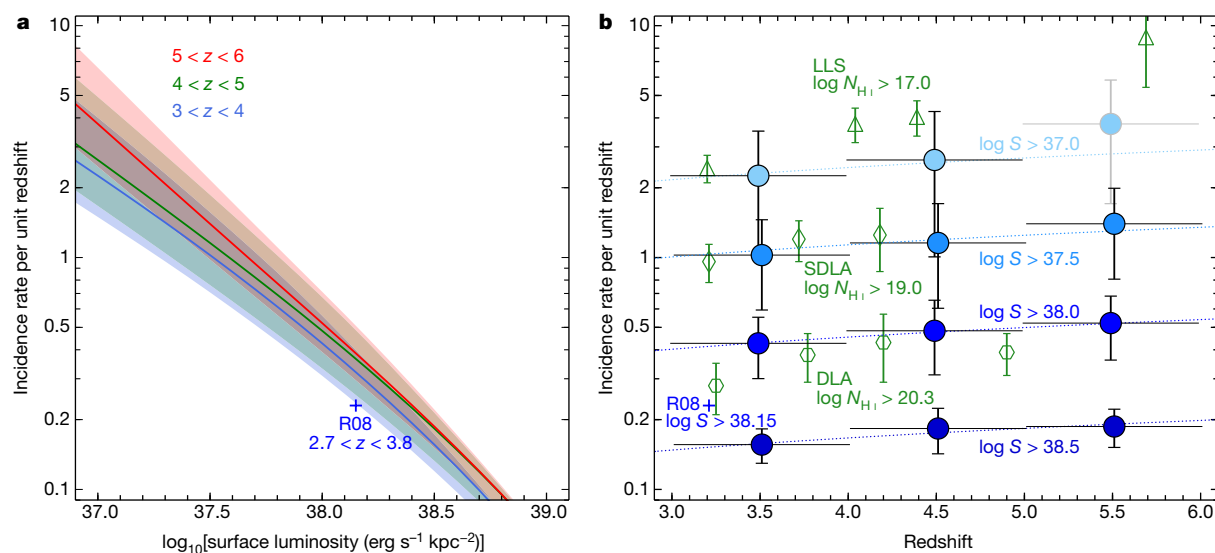


Fig. 4 | Incidence rates of Ly α emission and comparison with absorption measurements. **a**, Cumulative incidence rates of Ly α in emission, as a function of limiting surface luminosity; data are shown colour-coded for the three redshift ranges given. The shaded bands outline our estimated 1σ uncertainties. **b**, Blue filled circles show the inferred evolution of Ly α emission incidence rates with redshift, for four different thresholds in Ly α surface luminosity labelled by their $\log S$ values. In order to avoid overlapping error bars (1σ as defined in Methods) the symbols are offset by ± 0.01 in redshift. The top-right datapoint is plotted in light grey to indicate that it

involves extrapolation of the Ly α radial profiles and is less certain than the other points. The blue cross (R08) in both panels represents the only previous observational estimate of the Ly α emission incidence rate²⁰. Green open symbols show literature values of the cumulative incidence rates of atomic hydrogen measured from quasar absorption lines, for three commonly adopted limits in column density, $N_{\text{H I}}$ (triangles, Lyman limit systems¹², LLS; diamonds, sub-damped Ly α absorbers¹³, SDLA; hexagons, damped Ly α absorbers^{11,14}, DLA). The thin dotted lines show the expected trend for an intrinsically non-evolving population of Ly α emitters.

$\log_{10}[N(\text{H I}) (\text{cm}^{-2})] > 19$. The approximate equality of the incidence rates of H I absorbers with column densities $\log_{10}[N(\text{H I}) (\text{cm}^{-2})] > 19$ and Ly α -emitting regions with surface luminosities $\log_{10}(S_{\text{Ly}\alpha}) \gtrsim 37.5$ therefore suggests that in these regions we are observing the faint glow of ubiquitous circumgalactic atomic hydrogen. This result is robust with respect to modest deviations from the assumed azimuthal symmetry in the spatial distribution of the emitting gas (Methods). The main systematic uncertainty, not included in our error bars, lies in the fact that our sample is selected by its Ly α emission. We estimate that we captured roughly 50% of the galaxies at these redshifts (Methods), implying that the incidence rates of circumgalactic absorbers should be reduced by this factor before comparing with the emission incidence rates. These uncertainties do not affect our conclusion that most atomic hydrogen at redshifts 3 to 6 has now also been detected in emission.

Our results suggest that dn/dz_{em} increases mildly with redshift. In Fig. 4b we compare our measurements with the expected redshift dependence for an intrinsically non-evolving population of emitters. Such a population can be described by a constant incidence rate dn/dX_{em} per comoving path length X along the line of sight, a quantity commonly used in quasar absorption line studies. Within the redshift range considered here, a constant $dn/dX(z)$ translates into a $dn/dz(z)$ similar to the trend suggested by our data points. Because the Ly α luminosity function also shows very little evolution within the redshift range of our sample²⁸, we conclude that the properties of circumgalactic Ly α emission do not change much between redshifts 6 and 3. Nevertheless, given the error bars the data are also compatible with a constant $dn/dz_{\text{em}}(z)$, corresponding to a modest decrease of incidence rates per comoving path length.

What powers this low-level Ly α emission that we tentatively identify as originating from circumgalactic high-column-density absorbers? At large galactocentric distances, Ly α fluorescence of optically thick gas excited by the cosmic ultraviolet background (UVB) becomes a viable possibility³. The expected fluorescent Ly α surface brightness of an L LS at $z = 3$ has been calculated²⁹ and recently updated³⁰ for $z = 3.5$ as $s_{\text{Ly}\alpha, \text{UVB}} \approx 1.1 \times 10^{-20} \text{ ergs}^{-1} \text{ cm}^{-2} \text{ arcsec}^{-2}$, which is just about within the sensitivity range of our stacked data and consistent with the marginal signal at large radii in our lowest redshift bin. This suggests that at least some of the extremely faint Ly α emission detected by MUSE may be due to this omnipresent glow, opening a window to an important but previously invisible component of the cosmic matter distribution.

Online content

Any methods, additional references, Nature Research reporting summaries, source data, statements of data availability and associated accession codes are available at <https://doi.org/10.1038/s41586-018-0564-6>.

Received: 15 March 2018; Accepted: 1 August 2018;

Published online 1 October 2018.

1. Partridge, R. B. & Peebles, P. J. E. Are young galaxies visible? *Astrophys. J.* **147**, 868–886 (1967).
2. Hogan, C. J. & Weymann, R. J. Lyman-alpha emission from the Lyman-alpha forest. *Mon. Not. R. Astron. Soc.* **225**, 1P–5P (1987).
3. Gould, A. & Weinberg, D. H. Imaging the forest of Lyman limit systems. *Astrophys. J.* **468**, 462–468 (1996).
4. Wisotzki, L. et al. Extended Lyman α haloes around individual high-redshift galaxies revealed by MUSE. *Astron. Astrophys.* **587**, A98 (2016).
5. Leclercq, F. et al. The MUSE Hubble Ultra Deep Field Survey. VIII. Extended Lyman-alpha haloes around high- z star-forming galaxies. *Astron. Astrophys.* **608**, A8 (2017).
6. Hayes, M., Scarlata, C. & Siana, B. Central powering of the largest Lyman- α nebula is revealed by polarized radiation. *Nature* **476**, 304–307 (2011).
7. Cantalupo, S., Arrigoni-Battaia, F., Prochaska, J. X., Hennawi, J. F. & Madau, P. A cosmic web filament revealed in Lyman- α emission around a luminous high-redshift quasar. *Nature* **506**, 63–66 (2014).
8. Hennawi, J. F., Prochaska, J. X., Cantalupo, S. & Arrigoni-Battaia, F. Quasar quartet embedded in giant nebula reveals rare massive structure in distant universe. *Science* **348**, 779–783 (2015).
9. Steidel, C. C. et al. Diffuse Ly α emitting halos: a generic property of high-redshift star-forming galaxies. *Astrophys. J.* **736**, 160–177 (2011).
10. Momose, R. et al. Diffuse Ly α haloes around galaxies at $z = 2.2$ – 6.6 : implications for galaxy formation and cosmic reionization. *Mon. Not. R. Astron. Soc.* **442**, 110–120 (2014).

11. Péroux, C., McMahon, R. G., Storrie-Lombardi, L. J. & Irwin, M. J. The evolution of Ω_{HI} and the epoch of formation of damped Lyman α absorbers. *Mon. Not. R. Astron. Soc.* **346**, 1103–1115 (2003).
12. Songaila, A. & Cowie, L. L. The evolution of Lyman limit absorption systems to redshift six. *Astrophys. J.* **721**, 1448–1466 (2010).
13. Zafar, T. et al. The ESO UVES advanced data products quasar sample. *Astron. Astrophys.* **556**, A141 (2013).
14. Crighon, N. H. M. et al. The neutral hydrogen cosmological mass density at $z = 5$. *Mon. Not. R. Astron. Soc.* **452**, 217–234 (2015).
15. Bacon, R. et al. The MUSE second-generation VLT instrument. *Proc. SPIE* **7735**, 773508 (2010).
16. Bacon, R. et al. MUSE commissioning. *The Messenger* **157**, 13–16 (European Southern Observatory, Garching, 2014).
17. Bacon, R. et al. The MUSE 3D view of the Hubble Deep Field South. *Astron. Astrophys.* **575**, A75 (2015).
18. Casertano, S. et al. WFC2 observations of the Hubble Deep Field South. *Astron. J.* **120**, 2747–2824 (2000).
19. Beckwith, S. V. W. et al. The Hubble Ultra Deep Field. *Astron. J.* **132**, 1729–1755 (2006).
20. Rauch, M. et al. A population of faint extended line emitters and the host galaxies of optically thick QSO absorption systems. *Astrophys. J.* **681**, 856–880 (2008).
21. Laursen, P. & Sommer-Larsen, J. Ly α resonant scattering in young galaxies: predictions from cosmological simulations. *Astrophys. J.* **657**, L69–L72 (2007).
22. Barnes, L. A. & Haehnelt, M. G. Faint extended Ly α emission due to star formation at the centre of high column density QSO absorption systems. *Mon. Not. R. Astron. Soc.* **403**, 870–885 (2010).
23. Verhamme, A. et al. Lyman- α emission properties of simulated galaxies: interstellar medium structure and inclination effects. *Astron. Astrophys.* **546**, A111 (2012).
24. Haimes, Z., Spaans, M. & Quataert, E. Ly α cooling radiation from high-redshift halos. *Astrophys. J.* **537**, L5–L8 (2000).
25. Fardal, M. A. et al. Cooling radiation and the Ly α luminosity of forming galaxies. *Astrophys. J.* **562**, 605–617 (2001).
26. Furlanetto, S. R., Schaye, J., Springel, V. & Hernquist, L. Ly α emission from structure formation. *Astrophys. J.* **622**, 7–27 (2005).
27. Altay, G., Theuns, T., Schaye, J., Crighon, N. H. M. & Dalla Vecchia, C. Through thick and thin – H I absorption in cosmological simulations. *Astrophys. J.* **737**, L37 (2011).
28. Drake, A. B. et al. The MUSE Hubble Ultra Deep Field Survey. VI. The faint-end of the Ly α luminosity function at $2.91 < z < 6.64$ and implications for reionisation. *Astron. Astrophys.* **608**, A6 (2017).
29. Cantalupo, S., Porciani, C., Lilly, S. J. & Miniati, F. Fluorescent Ly α emission from the high-redshift intergalactic medium. *Astrophys. J.* **628**, 61–75 (2005).
30. Gallego, S. G. et al. Stacking the cosmic web in fluorescent Ly α emission with MUSE. *Mon. Not. R. Astron. Soc.* **475**, 3854–3869 (2018).

Acknowledgements We thank the ESO staff for the support that made these observations possible. L.W., J.K., R.S. and T.U. acknowledge support by the Competitive Fund of the Leibniz Association through grants SAW-2013-AIP-4 and SAW-2015-AIP-2. R.B., H.I., F.L. and M.A. are supported by the ERC advanced grant 339659-MUSICOS. J.B. acknowledges support by FCT grants UID/FIS/04434/2013 and IF/01654/2014/CP1215/CT0003 and by FEDER through COMPETE2020 (POCI-01-0145-FEDER-007672). J.R. acknowledges support from the ERC starting grant 336736-CALENDS. P.M.W. received support through BMBF Verbundforschung, grant 05A17BAA. T.C., N.B. and B.G. acknowledge support by ANR FOGHAR (ANR-13-BS05-0010-02). T.C. and N.B. were also supported by OCEVU Labex (ANR-11-LABX-0060), and by the A*MIDEX project (ANR-11-IDEX-0001-02) funded by the “Investissements d’avenir” French government programme. A.M.I. acknowledges support from MINECO through project AYA2015-68217-P. S.C. acknowledges support from Swiss National Science Foundation grant PP00P2_163824.

Reviewer information Nature thanks G. Micheva and the other anonymous reviewer(s) for their contribution to the peer review of this work.

Author contributions L.W. conceived the project. R.B. led the data acquisition and data reduction. R.B., J.B., E.C.H., H.I., J.K., K.B.S., T.U. and L.W. developed and performed the sample selection. L.W. analysed the data, with input from R.B., J.B. and P.M.W. S.C., P.R., J.S., M.S. and L.W. worked on the interpretation of the results. L.W. wrote the manuscript and produced the figures, with K.B.S. contributing to their design. All coauthors provided critical feedback to the text and helped shape the manuscript.

Competing interests The authors declare no competing interests.

Additional information

Extended data is available for this paper at <https://doi.org/10.1038/s41586-018-0564-6>.

Supplementary information is available for this paper at <https://doi.org/10.1038/s41586-018-0564-6>.

Reprints and permissions information is available at <http://www.nature.com/reprints>.

Correspondence and requests for materials should be addressed to L.W.

Publisher’s note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

METHODS

MUSE observations. MUSE is an integral field spectrograph mounted on Unit Telescope 4 of the ESO Very Large Telescope. In its Wide Field Mode it offers a $1' \times 1'$ field of view at a spatial sampling of $0.2'' \times 0.2''$, producing a datacube with 90,000 spatial pixels. Each spatial pixel contains a 475–935 nm spectrum with ~ 0.25 nm spectral resolution. The first of the two deep field observations used here was obtained in 2014, where we integrated for a total of 27 h on a single pointing in the Hubble Deep Field South¹⁸ (HDFS). The data reduction and construction of the first redshift catalogue in this field are summarized elsewhere¹⁷. The second deep MUSE exposure was obtained between 2014 and 2016 as part of a greater effort to perform a contiguous spectroscopic mapping of the Hubble Ultra-Deep Field¹⁹ (HUDF). These observations resulted in a $3' \times 3'$ mosaic with a mean integration time of 10 h, on top of which 21 h of additional exposure time were dedicated to a single MUSE pointing inside the HUDF. The characteristics of the MUSE-HUDF data set and the reduction process are described elsewhere³¹. Here we use only the two ultra-deep 1 arcmin^2 MUSE pointings, which for simplicity we refer to as HDFS and HUDF, respectively. The average spatial resolution (FWHM of the best-fitting Moffat³² function) in the combined coadded datacube, evaluated at 700 nm, is $0.66''$ for the HDFS and $0.63''$ for the HUDF.

The sample of Ly α emitters. Here we focus on galaxies marked by their Ly α emission (Ly α emitters, LAEs). Given the MUSE spectral range of 475–935 nm, LAEs can be detected over a redshift interval of $2.92 < z < 6.64$. In order to construct a homogeneous Ly α -selected sample, we ran our dedicated software LSDCat³³ (Line Source Detection and Cataloguing) with a signal-to-noise ratio threshold of 6 on both datacubes to produce a list of emission line objects, which we then inspected visually to assign redshifts. This resulted in a sample of 128 LAEs in the HDFS and 161 LAEs in the HUDF, respectively. Because of the strictly Ly α -based selection, these samples are not mere subsets of our previously published catalogues^{17,34} but also contain a few additional LAEs. The spatial distribution of the objects is visualized in Extended Data Fig. 1. For each LAE, the LSDCat software measures the spatial and spectral centroids and the integrated Ly α fluxes, quantities used in this study. Redshifts were assigned from the measured centroids of the Ly α emission line. While these centroids are known to be shifted with respect to the systemic redshifts by up to a few hundred km s^{-1} , accurate knowledge of the latter is not required for the present study (and was not available for our sample). Because of the decrease in instrument sensitivity close to the low and high wavelength cutoffs, and because of the crowding of OH night-sky emission lines towards the reddest wavelengths, we limited our sample to the redshift range $3 < z < 6$, which conveniently allowed us to define three broad redshift bins of $\Delta z = 1$ each. The final sample encompasses 119 LAEs in the HDFS and 151 LAEs in the HUDF, respectively. We provide the full sample as a data table (Supplementary Data) in machine-readable format containing the relevant quantities for this study: positions, redshifts, integrated Ly α fluxes, and flags indicating which objects were used for the direct projection and median stacking subsets. The newly found LAEs will also be included in a forthcoming public catalogue update for these fields.

Extraction of narrowband images. Each LAE enters into the current investigation as a pseudo-narrowband (NB) Ly α image, extracted from the datacube at the location of the three-dimensional Ly α centroid coordinates provided by LSDCat. These images were constructed in the following way: we first extracted a provisional Ly α spectrum from the continuum-subtracted datacube by summing over an unweighted circular aperture of radius $0.6''$. We then modelled the Ly α emission line profile as a Gaussian, which provided an improved line centroid as well as an approximate line width. Using this Gaussian approximation as spectral template, we performed a weighted summation of spectral layers of the datacube into a single NB image. While Ly α lines usually show some deviations from a single Gaussian, these deviations have negligible effect on the extraction results, except for cases of secondary line peaks ('blue bumps') which are not captured by our narrowband images. The maximum Gaussian FWHM for the extraction was set to 500 km s^{-1} in order to limit the noise. Compared to an unweighted summation over a given bandwidth this scheme provides a better signal-to-noise ratio, and the fractional weights ensure that the bandwidth always matches the actual line width, which matters especially for relatively narrow lines with $\text{FWHM} \lesssim 200 \text{ km s}^{-1}$. We verified that for broader lines the results are very similar to an unweighted summation. The blank-sky noise level in these NB images varies substantially, depending on the spectral bandwidth and on the wavelength of the Ly α line. A typical value of the pixel-to-pixel r.m.s. in regions with no detected emission is $5 \times 10^{-19} \text{ erg s}^{-1} \text{ cm}^{-2} \text{ pixel}^{-1}$, corresponding to a 1σ surface brightness limit of $10^{-19} \text{ erg s}^{-1} \text{ cm}^{-2} \text{ arcsec}^{-2}$ when averaged over an aperture of $1''$. This limit varies by a factor of ~ 2 between different objects.

The Ly α sky coverage from direct projection. We obtained a projected Ly α view of each field by coadding all extracted NB images, maintaining the position of each object in the plane of the sky. In order to reduce the noise in the coadded image, we introduced a truncation radius of $6''$ around the centroid of each LAE beyond which the NB data were set to zero for the coadding procedure. This was motivated

by the fact that beyond this radius we find generally no individually detectable Ly α emission around our objects. The projection therefore accounts only for the circumgalactic Ly α emission around detected sources. Any putative extended Ly α emission at the same redshift as a detected object, but outside the truncation radius, is ignored in the procedure. At the same time the truncation ensured that same-redshift pairs entered only once into the coadded image; if necessary we constructed additional masks by hand to ensure that this was always the case. Masks were also applied in a few cases of contamination by foreground emission lines or by continuum subtraction residuals from bright foreground objects such as stars. Finally, we removed 33 objects where Ly α falls close to a bright sky line causing significant sky-subtraction residuals in the narrowband images. Altogether the amount of budgeted Ly α flux in this approach should be seen as a strict lower limit to the true emission in the field.

With these provisions, we estimated the cumulative fractional sky coverage $f_{\text{Ly}\alpha}$ of Ly α emission brighter than $s_{\text{Ly}\alpha}$, documented in Extended Data Fig. 2. Without spatial filtering the surface brightness is given for a single pixel of only $0.2'' \times 0.2''$, resulting in extremely noisy images. We therefore filtered the images with a Gaussian of $\text{FWHM} = 7 \text{ pixels}$ ($1.4''$) which provided a good compromise between noise suppression, enforcing large-scale spatial coherence, and ensuring that flux redistribution from the central pixels into the outer parts can be neglected. The maximally reachable covering fraction is limited by random fluctuations due to noise, as can be seen very clearly in Extended Data Fig. 2: $f_{\text{Ly}\alpha}$ measured from the unfiltered data converges to considerably lower values than $f_{\text{Ly}\alpha}$ measured from the filtered data—which are of course also affected, but to a lesser degree. We have not attempted to push this approach any further, as for the main results of this paper we employed the stacking-and-insertion modelling approach described below. We did, however, perform a retrospective consistency check between the direct projection and the stacking approach, as follows: we took the idealized noise-free reconstructed model images obtained from the median-stacking analysis (the bottom-right image in Fig. 3a in the case of the HUDF), degraded them by adding realistic noise, and after spatial filtering we measured $f_{\text{Ly}\alpha}$ in the same way as in the real data. For the noise model we filled empty datacubes with normally distributed random numbers scaled to the effective noise in the actual data. From these noise-only cubes we extracted NB images with the same prescriptions as for the real LAEs, using the same spectral bandpasses and spatial masks, and coadded them to provide a random noise realization of the projected Ly α image, which we then added to the stacking-based model image. The grey curves in Extended Data Fig. 2 show that these very different approaches to measuring $f_{\text{Ly}\alpha}$ produce remarkably similar results, considering that the direct projection approach does not involve azimuthal averaging and is based on an incomplete sample without any completeness corrections. The only noteworthy discrepancy between the thick black and the thick grey lines in Extended Data Fig. 2 occurs around $\log_{10}(s_{\text{Ly}\alpha}) \approx -18.5$, mainly caused by the median-taking in the stacking process which removes the largest and most extended Ly α emitters. At these relatively high surface brightnesses, direct projection actually delivers a more realistic estimate of $f_{\text{Ly}\alpha}$ than the stacking approach.

Stacking analysis. We excised $20'' \times 20''$ MUSE-narrowband subimages centred on each source and put these into image stacks, separately for the three redshift intervals $3 < z < 4$, $4 < z < 5$ and $5 < z < 6$. Before the analysis, the data were subjected to a rigorous visual screening. In this step, 76 objects were removed from the stacks because their NB images were disturbed by sky subtraction residuals or residual emission from unrelated objects. 194 LAEs remained in the combined stacking sample. Bright foreground objects and the edges of the field of view were masked. We also ensured that when the NB image contained multiple objects at the same redshift, each spatial pixel contributed to the stack only once. We identified 13 double sources, 2 triples and one quadruple for which the subimage had thus to be divided up, by drawing a line through the image and assigning each pixel to only one object, or by using only pixels inside of a certain radius, typically $6''$. Finally, each stack was collapsed into a single image by computing the pixel-by-pixel median of all unmasked input image pixels. We chose the median instead of the mean to avoid the possibility of faint undetected companions enhancing the signal in the outskirts, and to make the stacked images robust against faint artefacts escaping from the visual screening. For each collapsed stack we also obtained a weight image containing in each pixel the number of input images that contributed to it after masking; these weight images are shown in Extended Data Fig. 3.

Profile extraction and error estimation. To detect the low-surface-brightness Ly α emission in the outer regions of our haloes, we extracted azimuthally averaged radial profiles from the median stacks (Fig. 2c), measured by averaging all pixels within each of a set of concentric annuli (Fig. 2b) defined as follows. Let r_i denote the outer radius of annulus i with $i = 1, \dots, 11$. For $i \leq 5$ we adopted constant annular widths, $r_i = i \times 0.2''$ (1 MUSE spatial pixel), for the outer annuli $i \geq 6$ we constructed a progression of increasing widths with the recursion formula $r_i = r_{i-1} + 10^{1.47(i-5)/7} \times 0.2''$. Thus, the last annulus has an outer radius $r_{11} = 9.97''$, just fitting into our $20'' \times 20''$ images and combining 4,660 MUSE spatial pixels into a single mean surface brightness measurement.

We estimated the uncertainties of these surface brightness profiles in two ways, by formal error propagation and empirically from empty regions in the data. The formal errors originate in the pixel variances of the MUSE databases, corrected for resampling effects³¹, propagated first into the NB images of individual objects and then into the median-combined stacks. For the latter step we used the property of the median that its variance is approximately $\pi/2$ times the variance of the mean. For the empirical noise calibration we shifted the whole LAE sample in wavelength while maintaining the spatial positions of all objects, thus defining empty regions by applying small redshift offsets. Offsetting in increments of ± 5 MUSE spectral pixels or ± 6.25 Å yielded 40 complete sets of as many empty regions as LAEs, which were then subjected to the same NB image extraction and median stacking procedure as the Ly α data. We estimated the noise from the dispersion of the extracted radial profiles between these 40 sets. The outcome of this experiment is presented in Extended Data Fig. 4. This figure thus provides the significance limits for the detection of very low-surface-brightness emission in our stacked data.

Although the propagated errors can be easily calculated independently of the object content of the databases, they do not account for systematics in the background subtraction. On the other hand, the empirically determined errors automatically include such systematics, but are subject to contamination of the supposedly empty regions by unrelated sources and by sky subtraction residuals, and hence they are likely to overestimate the true errors. In our experiment the median ratios between empirical and propagated errors for the outer profile regions ($r > 1''$) were (1.02, 1.33, 1.52) for the three redshift ranges. For this paper we conservatively adopted only the (larger) empirically determined errors for the profiles, and only these are shown in Fig. 2c and Extended Data Fig. 4. The empty regions experiment also revealed that the mean of the empty regions tends to become slightly negative at small radii $r < 2''$, implying that we may actually underestimate the surface brightnesses in the inner regions by a very small amount. Because this effect is at the $\sim 1\%$ level relative to the measured signal at these radii, we decided to neglect it.

Surface brightness modelling. We used GALFIT³⁵ to model the median-stacked images with smooth 2-dimensional Ly α surface brightness distributions. While in previous work^{4,5} we favoured a double exponential model with a compact core and an extended halo, that model was driven mainly by the high signal-to-noise ratio and high surface brightness regions within $\lesssim 10$ kpc. In the present study the emphasis is on the outer regions, where Fig. 2 suggests that the surface brightness distribution of the halo may show some flattening relative to a single exponential. Here we adopted a model consisting of a central point source plus a circular Sérsic³⁶ function, convolved with the point-spread function (PSF) at the relevant wavelengths, which provided good fits to all median-stacked images. To quantify the uncertainties in the fitted profiles we used the formal error estimates of GALFIT, but then increased these until the uncertainties of the fits were consistent with the empirically calibrated error bars of the directly extracted azimuthal profiles. These uncertainties are displayed as the shaded regions around the fitted profiles in Fig. 2. Extended Data Table 1 provides the numerical values of the fit parameters and their uncertainty estimates.

We make the central assumption that the shapes of the Ly α haloes of LAEs do not, in the statistical average, depend on their luminosities. While in ref. ⁵ we found no evidence for such a dependence, most of those objects have Ly α -luminosities $L_{\text{Ly}\alpha} > 10^{42} \text{ erg s}^{-1}$, whereas our current sample has a median $\log_{10}[L_{\text{Ly}\alpha} (\text{erg s}^{-1})]$ of only 41.7. Here we use the median stacks to probe deeper into the validity of the above assumption. In Extended Data Fig. 5 we present a comparison between the radial profiles extracted from the median stack of the full sample and of a subset with Ly α luminosities $L_{\text{Ly}\alpha} > 10^{42} \text{ erg s}^{-1}$. There are (18, 13, 6) objects at $z = (3-4, 4-5, 5-6)$ meeting this criterion. The median-stack profiles of the luminous subsets resemble those of the full sample remarkably well. This comparison supports the validity of our self-similarity approximation across the luminosity range of our sample. We plan to revisit this point and related aspects in a follow-up study of Ly α halo profile shapes in stacked MUSE data.

We used the analytic profile fits as templates to construct idealized representations of all LAEs in the two fields, including those objects previously removed from the stacking subsample. All model LAEs at a given redshift range were assigned to have the same spatial surface brightness distribution, but rescaled to the actually observed Ly α fluxes of each real object. Furthermore, the GALFIT models were approximately corrected for PSF blurring by using a delta function (in fact a very narrow Gaussian) as PSF when reconstructing the two-dimensional model templates. Since for each object the template was rescaled to match the measured Ly α flux, this implies that the modelling of the brightest LAEs involved a certain degree of extrapolation of the profiles. We demonstrate below that the contribution of extrapolated emission to the incidence rates is small (see also Extended Data Fig. 7).

Determination of incidence rates. We estimated the Ly α emission incidence rates directly from the LAE samples as the sum of circular cross-sections $\pi r_{\text{iso},i}^2$, where $r_{\text{iso},i}(S_{\text{Ly}\alpha})$ is the isophotal extent of the model of object i at a given surface brightness level. The incidence rate is then

$$\frac{dn}{dz}(S_{\text{Ly}\alpha}) \approx \frac{1}{A_{\text{FoV}}} \sum \frac{n_{\text{obj}}}{\Delta z_i} \sum_{i=1}^{n_{\text{obj}}} \pi r_{\text{iso},i}^2(S_{\text{Ly}\alpha}) \quad (1)$$

where A_{FoV} is the area of the field of view, Δz_i is the redshift path length over which object i would be part of the flux-limited sample, and where the summation is carried out over all objects in the redshift range. The normalization of dn/dz to a quantity per unit redshift was in our case conveniently provided by the widths of our adopted redshift intervals.

Since $S_{\text{Ly}\alpha}$ decreases with increasing redshift as $(1+z)^4$, we decided to move to distance-independent surface luminosities $S_{\text{Ly}\alpha}$. While this quantity is not much used in the literature, we prefer working with intrinsic object properties over rescaling observed quantities to some fiducial reference redshift. The transformation is $\log_{10}(S_{\text{Ly}\alpha}) = \log_{10}(S_{\text{Ly}\alpha}) + 4\log_{10}(1+z) + 54.71$ when both s and S are given in cgs units. The right-hand ordinate of Extended Data Fig. 5 provides a quick-look visual calibration of the conversion.

Correction for sample incompleteness. Our LAE sample certainly suffers from incompleteness close to the flux limit, with faint objects getting selected only if their Ly α emission is sufficiently point-like. Selection effects for the detection of LAEs in the MUSE-HUDF survey have been investigated in detail²⁸, and it was shown that the transition from 80% to 20% detection probability extends over ~ 0.5 dex in line flux. The Ly α incidence rates calculated from equation (1) are therefore biased low, missing the contributions from undetected but presumably existing objects. In order to correct for this incompleteness, we replaced the summation over the observed sample by an integration over the full survey volume, assuming that the intrinsic distribution of Ly α luminosities follows the luminosity function determined in ref. ²⁸. Since the self-similarity approximation of the extended Ly α emission implies a unique relation between the total flux $F_{\text{Ly}\alpha}$ of an object and its isophotal radius, $r_{\text{iso}} = r_{\text{iso}}(F_{\text{Ly}\alpha}, S_{\text{Ly}\alpha})$, we can predict the emission cross-sections from only the Ly α luminosities. The total incidence rate dn/dz for a given redshift range (z_1, z_2) follows as

$$\frac{dn}{dz}(S_{\text{Ly}\alpha}) = \frac{1}{(z_2 - z_1)A_{\text{FoV}}} \int_{z_1}^{z_2} \int_{\ell_{\min}}^{\ell_{\max}} \pi r_{\text{iso}}^2[S_{\text{Ly}\alpha}, F_{\text{Ly}\alpha}(\ell, z)] \phi(\ell, z) d\ell dV_c(z) \quad (2)$$

where $dV_c(z)$ is the differential comoving volume element at redshift z , $\ell \equiv \log_{10}(L_{\text{Ly}\alpha}) = \log_{10}(4\pi d_L^2 F_{\text{Ly}\alpha})$ (in units of erg s^{-1} , with the cosmological luminosity distance d_L), and $\phi(\ell, z)d\ell$ is the differential Ly α luminosity function at redshift z . We parameterized $\phi(\ell)$ as a redshift-independent Schechter function with $\ell^* = 42.59$, $\phi^* = 2.138 \times 10^{-3} \text{ Mpc}^{-3}$, and $\alpha = -1.93$, which is a good overall fit to the completeness-corrected LAE sample²⁸. While the upper luminosity integration limit ℓ_{\max} does not matter much as $\phi(\ell)$ approaches zero very quickly for increasing ℓ , choosing a value for the lower integration limit ℓ_{\min} is less straightforward: Although faint LAEs have small isophotal radii, they are also numerous and thus contribute non-negligibly to the integrated cross-section. The integral converges only for $\ell_{\min} \lesssim 40.5$, but at the expense of including large numbers of hypothetical ultrafaint LAEs into the budget that are well below the current detection limits. As our 'best guess' we adopted $\ell_{\min} = 41.0$, which is slightly brighter than the faintest detected LAEs in our actual sample. Extended Data Fig. 6 provides a synopsis of the different approaches to estimate dn/dz from our data. These plots also show that the magnitude of the completeness correction is by far the dominant source of uncertainty for dn/dz . We therefore adopted as the lowest reasonable limit the values of dn/dz without any completeness correction (that is, from equation (1)) obtained from the somewhat 'emptier' HDFs. As an upper limit we took the asymptotic result from integrating equation (2) with $\ell_{\min} = 40.0$. For the presentation in Fig. 4 we interpreted these lower and upper bounds as $\pm 2\sigma$ limits, but plotted only the 1σ error envelopes in accordance with the usual conventions.

Discussion of systematic errors. We first consider to what extent the integrated values of dn/dz depend on extrapolations of the rescaled Ly α profiles beyond the radial range over which they were constructed. To address this question, we constructed the cumulative distribution of the contributions to the total incidence rate sum or integral as a function of their isophotal radii. The results of the completeness-corrected integration of dn/dz (equation (2)) with $\ell_{\min} = 41.0$ (our 'best guess' approach) are shown in Extended Data Fig. 7. These plots demonstrate that at all redshifts and for all levels of $S_{\text{Ly}\alpha}$ except the lowest, more than 80% of the total incidence rates is contributed by emission from $r_{\text{iso}} < 5''$. Only for $\log_{10}[S_{\text{Ly}\alpha, \text{lim}} (\text{erg s}^{-1} \text{ kpc}^{-2})] = 37$ and $5 < z < 6$ there is a substantial extrapolated contribution, which is why we consider this point as uncertain and plotted it in light grey in Fig. 4b.

A strong assumption made in our analysis is the azimuthal symmetry of the Ly α emission, which is certainly not strictly correct. We now quantify the biases arising from the circularization of a non-axisymmetric signal through median stacking. Extended Data Fig. 8a–c shows how an elliptical two-dimensional surface brightness distribution and the corresponding isophotal cross-sections are

modified when the cross-sections are estimated from an azimuthally averaged profile. The circularized profile is broadened and the cross-sections at given surface brightnesses are overestimated, but the effect is quite small. More relevant for our analysis, Extended Data Fig. 8d–f demonstrates that median stacking of randomly oriented elongated objects delivers isophotal cross-sections that are systematically smaller than the true values. While the above calculations are based on a rather simple source model, the conclusion can be qualitatively generalized to other non-axisymmetric surface brightness distributions. The median stacking ensures that the derived emission cross-sections, and consequently the inferred Ly α sky coverage and incidence rates, are rather under- than overestimated.

While the small scale distribution of Ly α emission remains unknown at these surface brightness levels, we have some idea of the projected covering fraction f_{HI} of neutral hydrogen close to galaxies. Ref. ³⁷ used the cosmological EAGLE simulation³⁸ to show that f_{HI} depends on many parameters: distance to galaxy centre, column density, redshift, halo mass, environment. Measurements³⁹ of the H I absorption line close to galaxies at $z \approx 2.5$ give $f_{\text{HI}} = 0.3 \pm 0.14$ for Lyman limit systems within one virial radius (r_{vir}); there are so far no good measurements for $z > 3$. Simulations predict that f_{HI} increases rapidly towards higher redshift³⁷, and we expect $f_{\text{HI}} \approx 0.4$ – 0.8 for $r < r_{\text{vir}}$ at the redshifts of our sample. The virial radii of our LAEs are unknown, but are predicted to be around 30 kpc or less⁴⁰. Consulting again Extended Data Fig. 7, we see that except for $\log_{10}[S_{\text{Ly}\alpha, \text{lim}} (\text{erg s}^{-1} \text{kpc}^{-2})] = 37$ and $5 < z < 6$, more than 80% of the integrated Ly α emission incidence rates come from radii less than 30 kpc, that is, from within one virial radius. Unless the Ly α -emitting gas has a very different spatial distribution from the general circumgalactic H I, the covering fractions f_{HI} are expected to be sufficiently close to unity that the systematic errors from any non-axisymmetry of our derived incidence rates should be small.

A rather different systematic error arises from the limitation of our sample to galaxies selected by their Ly α emission. If not all galaxies are LAEs, then there is circumgalactic H I gas contributing to high-column-density absorption systems which is not included in our budget of extended Ly α emission. The fraction of galaxies at $z > 3$ showing detectable Ly α emission depends strongly on the selection criteria. While only some 10%–20% of continuum-bright galaxies at these redshifts are⁴¹ strong LAEs with Ly α rest-frame equivalent widths (EW) greater than 50 Å, this fraction probably increases to $\sim 50\%$ if weaker emitters are also included^{9,42}. There are indications that the fraction of strong emitters may even be considerably larger than 50% for very low luminosity galaxies and/or higher redshifts^{41,43}. It seems thus plausible to estimate that our Ly α selection captured roughly half of all galaxies at these redshifts. If the other 50% have a circumgalactic medium similar to that of the LAEs except for the lack of Ly α emission (this is a very uncertain assumption), the incidence rates of the circumgalactic absorbers in Fig. 4 should be shifted downwards by ~ 0.3 dex to account only for the LAE fraction. Interpolating between the measured values, a surface luminosity level of $\log_{10}[S_{\text{Ly}\alpha} (\text{erg s}^{-1} \text{kpc}^{-2})] \approx 37.5$ would then have roughly the same incidence rate as absorbers with $\log_{10}[N(\text{H I}) (\text{cm}^{-2})] \approx 18$, less than the limit for SDLAs but still

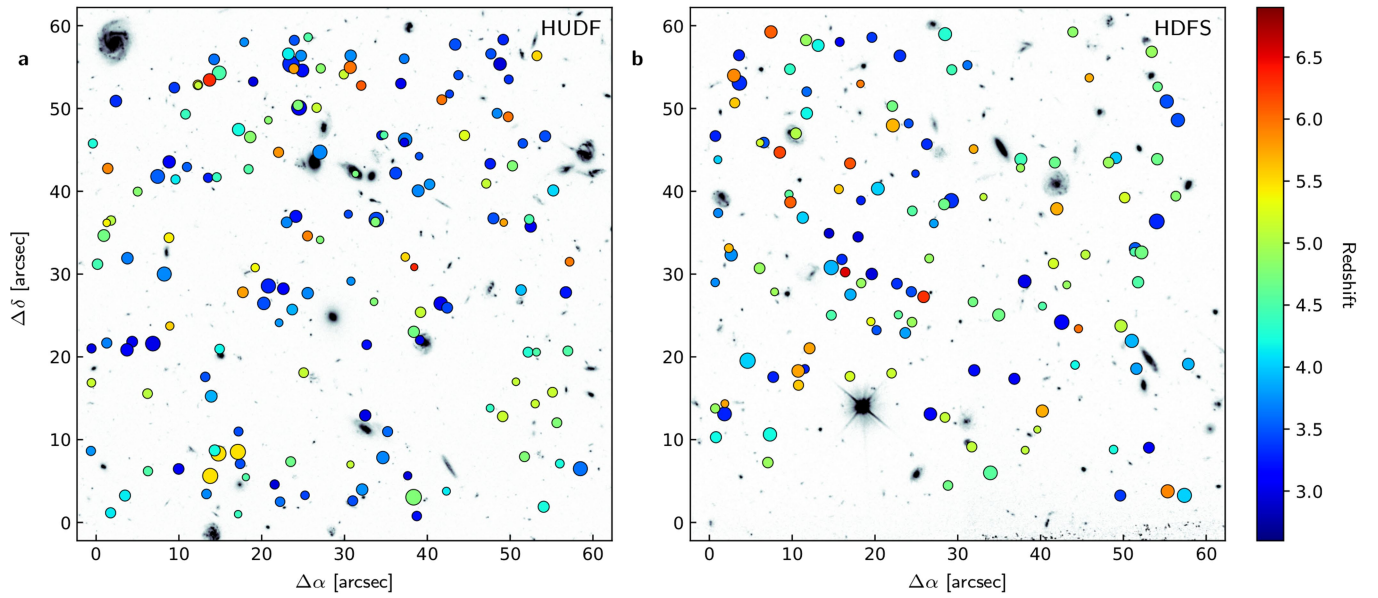
optically thick to Lyman continuum radiation. On the other hand, there may also be non-LAE galaxies with still undetected faint Ly α haloes, similar to those found in ref. ⁹, which would increase the Ly α incidence rates even further.

Code availability. This study was carried out with several small custom routines, mostly in Python. While the analysis steps are described in the paper, individual pieces of code can be provided upon request.

Data availability


The observations of the HUDF discussed in this paper were made using European Southern Observatory (ESO) Telescopes at the La Silla Paranal Observatory under programme IDs 094.A-0289, 095.A-0010, 096.A-0045 and 096.A-0045. The corresponding data are available on the ESO archive at <http://archive.eso.org/cms.html>. The data of the HDFS were obtained during MUSE commissioning observations and are available at <http://muse-vlt.eu/science/hdfs-v1-0/>.

31. Bacon, R. et al. The MUSE Hubble Ultra Deep Field Survey. I. Survey description, data reduction, and source detection. *Astron. Astrophys.* **608**, A1 (2017).
32. Moffat, A. F. J. A theoretical investigation of focal stellar images in the photographic emulsion and application to photographic photometry. *Astron. Astrophys.* **3**, 455–461 (1969).
33. Herenz, E. C. & Wisotzki, L. LSDCat: Detection and cataloguing of emission-line sources in integral-field spectroscopy datacubes. *Astron. Astrophys.* **602**, A111 (2017).
34. Inami, H. et al. The MUSE Hubble Ultra Deep Field Survey. II. Spectroscopic redshifts and comparisons to color selections of high-redshift galaxies. *Astron. Astrophys.* **608**, A2 (2017).
35. Peng, C. Y., Ho, L. C., Impey, C. D. & Rix, H.-W. Detailed structural decomposition of galaxy images. *Astron. J.* **124**, 266–293 (2002).
36. Sersic, J. L. *Atlas de Galaxias Australes* Vol. 1 (Observatorio Astronomico, Cordoba, 1968).
37. Rahmati, A. et al. The distribution of neutral hydrogen around high-redshift galaxies and quasars in the EAGLE simulation. *Mon. Not. R. Astron. Soc.* **452**, 2034–2056 (2015).
38. Schaye, J. et al. The EAGLE project: simulating the evolution and assembly of galaxies and their environments. *Mon. Not. R. Astron. Soc.* **446**, 521–554 (2015).
39. Rudie, G. C. et al. The gaseous environment of high- z galaxies: precision measurements of neutral hydrogen in the circumgalactic medium of $z \sim 2$ – 3 galaxies in the Keck Baryonic Structure Survey. *Astrophys. J.* **750**, 67 (2012).
40. Garel, T. et al. The UV, Lyman α , and dark matter halo properties of high-redshift galaxies. *Mon. Not. R. Astron. Soc.* **450**, 1279–1294 (2015).
41. Stark, D. P., Ellis, R. S., Chiu, K., Ouchi, M. & Bunker, A. Keck spectroscopy of faint $3 < z < 7$ Lyman break galaxies – I. New constraints on cosmic reionization from the luminosity and redshift-dependent fraction of Lyman α emission. *Mon. Not. R. Astron. Soc.* **408**, 1628–1648 (2010).
42. Caruana, J. et al. The MUSE-Wide survey: a measurement of the Ly α emitting fraction among $z > 3$ galaxies. *Mon. Not. R. Astron. Soc.* **473**, 30–37 (2018).
43. Stark, D. P., Ellis, R. S. & Ouchi, M. Keck spectroscopy of faint $3 < z < 7$ Lyman break galaxies: a high fraction of line emitters at redshift six. *Astrophys. J.* **728**, L2 (2011).



Extended Data Fig. 1 | Spatial distribution and redshifts of the Ly α emitter sample. **a**, The region observed with MUSE in the Hubble Ultra Deep Field (HUDF), **b**, the same for the Hubble Deep Field South (HDFS). Each Ly α emitter is represented by a circle colour-coded by redshift (key at right) and with a radius scaled by the integrated Ly α flux of the object.


α , right ascension; δ , declination. The objects shown here constitute the full sample. There are several cases of significant crowding of unequal-redshift objects separated by less than a few arcseconds in projection. The underlying greyscale images show the two fields as seen with the HST.



Extended Data Fig. 2 | Ly α sky coverage from direct projection.


a–c, Data from HUDF; **d–f**, data from HDF5. Greyscale images display the projected and coadded Ly α emission over the redshift range $3 < z < 6$ separately for the two observed fields: **a–d**, without any spatial filtering; **b–e**, after Gaussian filtering with $\text{FWHM} = 1.4''$. The image **b** is identical to the blue overlay in Fig. 1. The light brown areas delineate the MUSE field of view and indicate masked bright foreground objects. **c–f**, the

resulting fractional sky coverage; the black dashed line representing the unfiltered images and the thick black solid line representing the spatially filtered images. The thin grey line shows the result from the stacking analysis (Fig. 3) for comparison, and the thick grey line is a fiducial reconstruction of the sky coverage derived from a noisy and filtered version of the stacking-based model images. See the discussion in Methods for an interpretation of these figures.



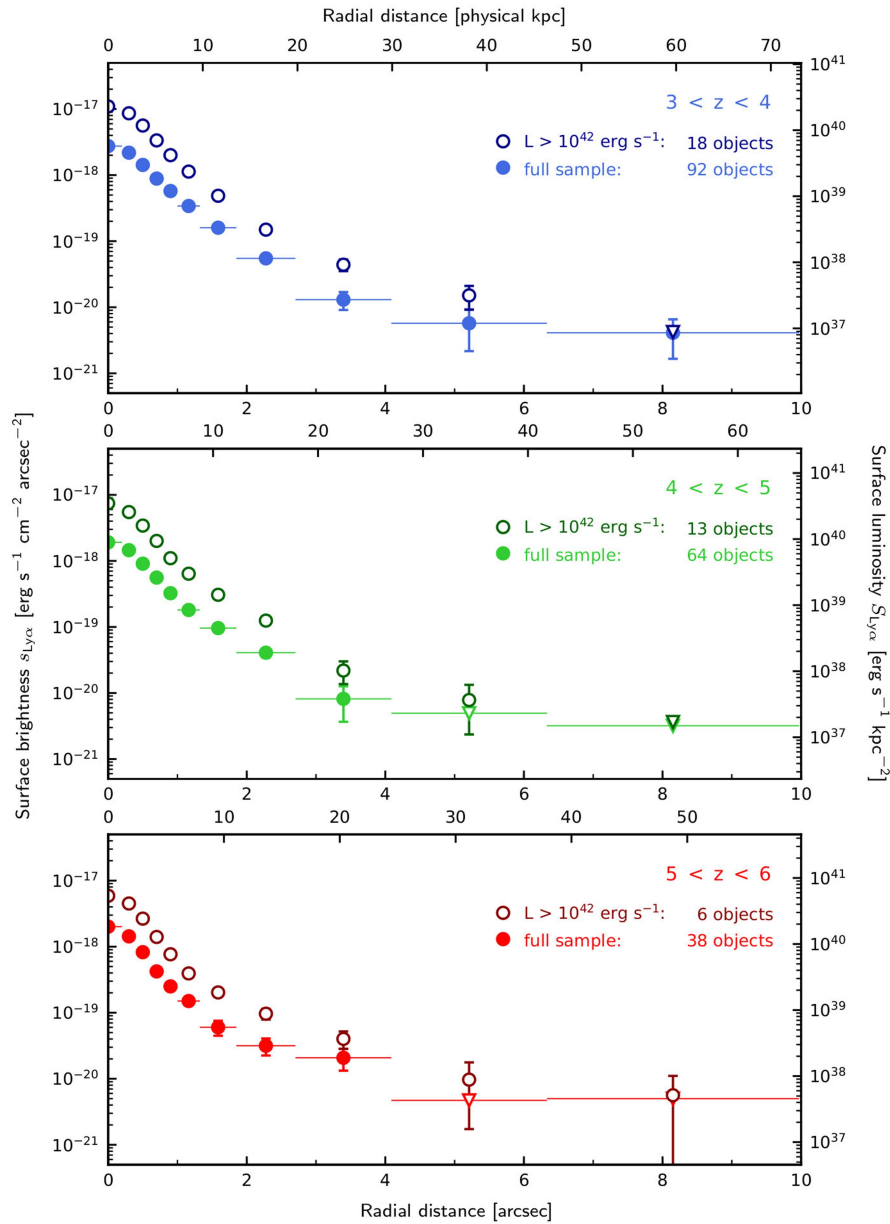
Extended Data Fig. 3 | Number of images contributing to the median stacks. a–c, Data for redshift ranges $3 < z < 4$ (a), $4 < z < 5$ (b) and $5 < z < 6$ (c). The colour code represents, for each pixel in a median-stacked

image, the number of original image pixels contributing to it. This number differs from the total number of objects in a given stack because of the masks applied to several of the contributing images.




Extended Data Fig. 4 | Error calibration and sensitivity of azimuthally averaged profiles. Shown are ‘zoom views’ into the profiles (surface brightness versus radial distance) at very low surface brightnesses, with linear ordinate scaling so that negative measurements can be displayed. Data for three redshift ranges are shown: top, $3 < z < 4$; middle, $4 < z < 5$; and bottom, $5 < z < 6$. The small open circles represent the surface

brightnesses measured in 40 realizations of a median-stacking analysis of empty regions as described in Methods. The filled symbols reproduce the data points from Fig. 2c. The horizontal bars again indicate the widths of the annuli, and the vertical error bars are based on the 1σ scatter of the empty field median-stack profiles; these were also adopted as error bars for the data points in Fig. 2c.




Extended Data Fig. 5 | A test of the self-similarity assumption for different Ly α luminosities. a–c, Comparison of azimuthally averaged radial profiles of median-stacked Ly α images above a minimum Ly α luminosity L (open circles) and with no such cut ('full sample', filled circles), for three redshift ranges (top, $3 < z < 4$; middle, $4 < z < 5$; and bottom, $5 < z < 6$). As in Fig. 2, the vertical bars on the data points

quantify the 1σ surface brightness measurement errors, while the horizontal bars (drawn only for the filled symbols) indicate the widths of the annuli. Inverted triangles indicate upper limits. The right-hand ordinate provides the conversion from apparent surface brightnesses to redshift-corrected surface luminosities, evaluated at the central redshift of each bin.



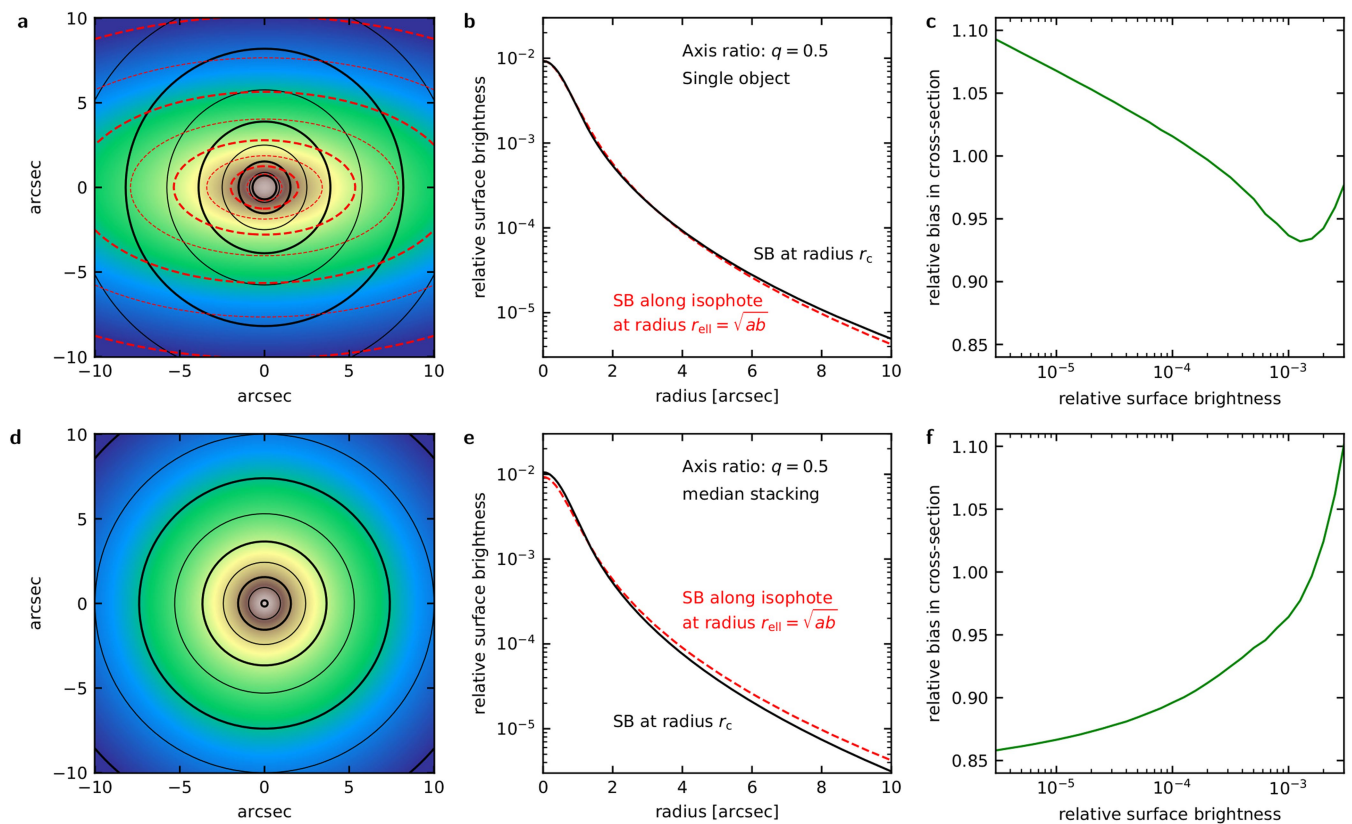
Extended Data Fig. 6 | Comparison of approaches to determine the Ly α emission incidence rates. Each panel shows the cumulative incidence rate as a function of limiting surface luminosity for the specified redshift range (left, $3 < z < 4$; middle, $4 < z < 5$; and right, $5 < z < 6$), estimated by different methods: direct summation of Ly α cross-sections over the sample without correcting for incompleteness (equation (1) in the Methods section, thin lines), and integrating over the completeness-corrected luminosity function following equation (2) in the Methods section, using lower integration limits of $\ell_{\min} = 41.0$ (best guess, thick

solid line) and $\ell_{\min} = 40.0$ (asymptotic case, dashed line), respectively. The shaded error bands for direct summation are dominated by field-to-field variance between the HUDF and the HDFs, with the upper envelope tracing the HUDF and the lower envelope tracing the HDFs results. For the luminosity function integration the error bands on these curves incorporate only the statistical uncertainties of the median-stacked profiles. The two thick dotted lines indicate the finally adopted lower and upper 2σ bounds on the ‘best guess’ results shown in Fig. 4.



Extended Data Fig. 7 | Cumulative contributions to incidence rates for different radii. Each panel (left, $3 < z < 4$; middle, $4 < z < 5$; and right, $5 < z < 6$) shows the cumulative fractional contributions of objects with different isophotal radii to the integrated Ly α emission cross-section

dn/dz , using equation (2) with $\ell_{\min} = 41$. The four lines in each panel represent, from left to right and with decreasing line width, surface luminosity limits $\log_{10}[S_{\text{Ly}\alpha} (\text{erg s}^{-1} \text{kpc}^{-2})] = 38.5, 38, 37.5, 37$.



Extended Data Fig. 8 | Bias of estimated cross-sections if the emission is non-axisymmetric. **a**, Model image of an elongated surface brightness distribution, normalized to an integrated flux of 1, following an elliptical Sersic law with axis ratio $q = 0.5$ and smoothed with a Gaussian of $0.8''$ FWHM. The colour code represents relative surface brightness (SB), and the red-dashed contours trace the isophotes at 0.5 dex separation. The black circles represent the radii where an azimuthally averaged profile over circular annuli gives the same surface SB values as the corresponding isophote. **b**, Radial profiles of the model image. The red-dashed line represents the input SB law as a function of generalized radius $r_{\text{ell}} = \sqrt{ab}$ where a, b are the major and minor axes of an isophote. The black line shows the profile obtained from azimuthal averaging over circular annuli

against radius r_c . **c**, Ratios between the true isophotal cross-sections πab and those estimated from the circularized profile as πr_c^2 (that is, the ratios of the areas of the black circles and the corresponding red-dashed ellipses in **b**), as a function of surface brightness. **d**, Median-stacked image of an ensemble of 180 model objects with properties each as in **a**, but rotated in position angle between 0° and 180° in steps of 1° . The colour code again represents relative SB, and the black circles show the resulting isophotes at 0.5 dex separation. **e**, The black line traces the radial profile of the median-stacked image in **d**. The red-dashed line is the true elliptical SB distribution of a single object (same as in **b**). **f**, Ratios of cross-sections obtained from the median stack to the true isophotal ones in a single image, as a function of surface brightness.

Extended Data Table 1 | Values of the best-fit parameters for the analytic profiles

z range	F_{h}	$r_{\text{eff,h}}$	n_{h}	F_{ps}	FWHM_{PSF}
3–4	1488 ± 83	0.86 ± 0.11	2.8 ± 1.1	232 ± 50	0.703
4–5	931 ± 82	0.90 ± 0.18	3.3 ± 1.9	150 ± 40	0.654
5–6	1002 ± 164	1.67 ± 0.86	6.5 ± 5.1	167 ± 34	0.606

These values were obtained by applying GALFIT to the median-stacked images. For each of the three adopted redshift ranges (column 1), the first three parameters characterize the circular Sersic model used to describe the Ly α haloes: halo flux F_{h} (in 10^{-20} erg s $^{-1}$ cm $^{-2}$), effective radius $r_{\text{eff,h}}$ (in arcsec), and Sersic index n_{h} (dimensionless), followed by the flux of the point-like component F_{ps} (same units as F_{h}). The quoted errors are 1σ uncertainty estimates. The last column provides the seeing (FWHM of the mean PSF in arcsec) at the appropriate wavelengths.